

HP Insight Cluster Management Utility



Efficient, robust hyperscale lifecycle management

HP Insight Cluster Management Utility (HP Insight CMU) simplifies the management and the rapid deployment of large Linux clusters and groups of clusters. HP Insight CMU's flexibility and efficient interface enable optimal workload performance and reduced administrative costs.

HP Insight Cluster Management Utility

A simple and affordable cluster toolkit, HP Insight CMU is used to configure and install cluster operating environments, monitor cluster and node metrics, and remotely manage resources. HP Insight CMU serves as a powerful tool for installing Linux software images, including middleware such as Message Passing Interface (MPI), and job schedulers. You can use HP Insight CMU anywhere you need to manage a number of standalone systems that are similar in hardware and software configuration.

Designed to manage large clusters, such as those found in high-performance computing (HPC) environments, HP Insight CMU operates independently of the underlying operating system, interconnects, and software applications. This flexible design enables HP Insight CMU to operate across multiple customized software stacks, and manage groups of nodes running in parallel with different operating system distributions.

From an historical perspective

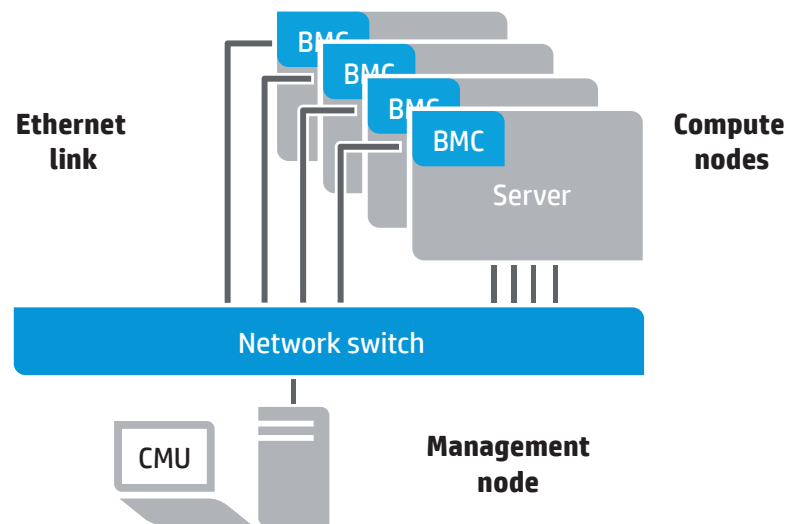
There is an enormous demand for industry-standard clusters. This demand was created by the emergence of Linux and Windows® as the most popular HPC operating systems, the rise in performance of x84-64 processors with Intel® Xeon®, AMD Opteron, and other chips, and Gigabit Ethernet and InfiniBand.

While the HPC cluster architecture has proven its efficiency in terms of performance and price, managing configurations of hundreds or thousands of systems is challenging. To support a range of customers—each with unique cluster implementations—HP Insight CMU was designed to manage key tasks, independent of operating system variant, MPI, or other software components. HP Insight CMU is an ideal tool for customers implementing a cluster solution utilizing multiple open source and/or third-party software components, and customers who need a simple interface for installation and administration of HPC clusters.

HP Insight CMU was first developed in the late 1990s. It now supports HP ProLiant servers—including the latest HP ProLiant SL series, HP Moonshot servers, and HP BladeSystem infrastructures—and includes recent enhancements to enable accelerator-based computation. Companies around the world, including many organizations on the Top500¹ list, use HP Insight CMU to manage a range of clusters.

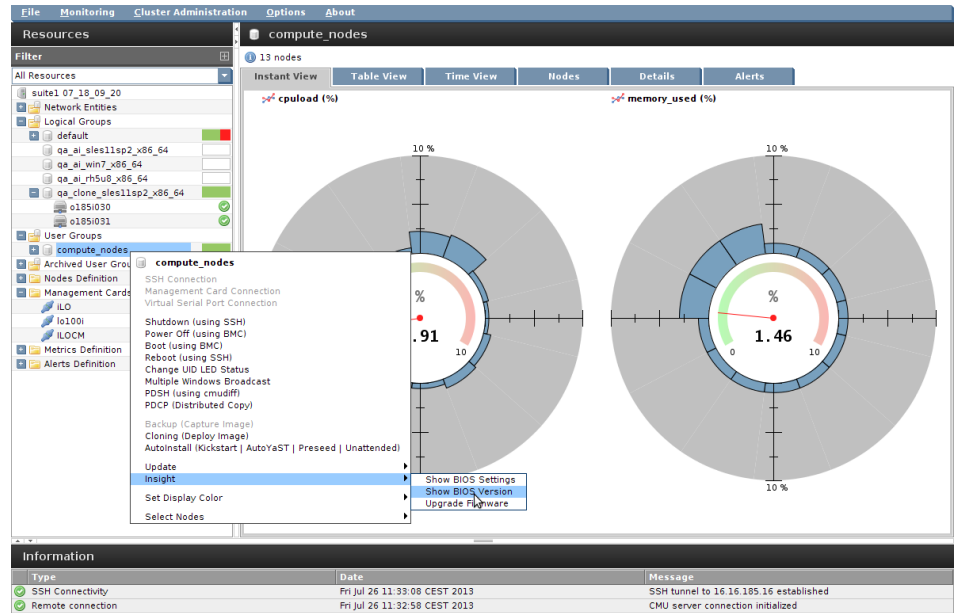
HP Insight CMU is supported on Red Hat Linux, SUSE Linux Enterprise Server (SLES), and multiple community distributions. In addition, HP Insight CMU has been augmented to enable Windows provisioning of HP Moonshot servers. HP Insight CMU has an application interface designed for easy extensibility of the scalable framework. This interface has been used to integrate several popular HPC management products with HP Insight CMU, including Altair Engineering PBS Professional, Adaptive Computing Moab, and Platform LSF.

Figure 1. GUI sample screenshot displaying global monitoring information of a cluster



¹Semiannual listing of the Top 500 performing systems, as listed at Top500.org

Figure 2. GUI sample screenshot displaying global monitoring information of a cluster



HP Insight CMU architecture

HP Insight CMU manages a cluster from a single management or “head” node. With DHCP, NFS, and tftp services, HP Insight CMU can PXE-boot the cluster nodes during installation and any subsequent software updates. Using features of the baseboard management controller (BMC) card of each compute node server (for example, Integrated Lights-Out [iLO] or Lights-Out 100), HP Insight CMU enables:

- Remote text-based console availability during all server states (setup, boot, OS, or halted)
- Remote control of server power, regardless of the state of the server (even if the server is off)
- Remote BIOS setup

HP Insight CMU features

Designed to streamline management of many thousands of compute nodes, HP Insight CMU ships with a complete Java GUI for remote management of the cluster, and a CLI for simple interactive and/or scriptable cluster management from the head node.

HP Insight CMU provides central access to all nodes in the cluster in several ways—from individual access to the OS or BMC of a node to a parallel distributed shell over all nodes. The GUI provides a multi-window broadcast capability that displays OS or serial console activity from a selected group of nodes, and broadcasts commands from a single keyboard session.

With a few mouse clicks in the GUI or a single command in the CLI, HP Insight CMU can shut down, boot-up, reboot, or power-off any selection of nodes. Other features available from the default dropdown menu are automated firmware update, BIOS consistency verification, and BIOS setting modification.

You can use HP Insight CMU to deploy software to the cluster hardware in several ways—the most efficient and scalable method is the HP Insight CMU backup and cloning mechanism. You configure one compute node, and HP Insight CMU will copy and distribute that disk image to some or all cluster nodes using an intelligent hierarchical process. HP Insight CMU can store multiple disk images, which means you can deploy different software stacks within a single cluster. HP Insight CMU also includes tools to support installation of GPU drivers and CUDA software.

Table 1. HP Insight CMU support matrix

	Processor	Linux distribution	Cluster nodes supported	Cluster interconnects
HP Cluster Platform 3000	Intel Xeon	<ul style="list-style-type: none"> • Red Hat • SUSE 	<ul style="list-style-type: none"> • ProLiant DL160, DL380 • Blades BL460c 	<ul style="list-style-type: none"> • InfiniBand • Gigabit Ethernet
HP Cluster Platform 4000	AMD Opteron	<ul style="list-style-type: none"> • Red Hat • SUSE 	<ul style="list-style-type: none"> • ProLiant DL165, DL385, DL585 • Blades BL465c 	<ul style="list-style-type: none"> • InfiniBand • Gigabit Ethernet
HP Cluster Platform 3000SL/4000SL	Intel Xeon	<ul style="list-style-type: none"> • Red Hat • SUSE 	<ul style="list-style-type: none"> • ProLiant SL230x, SL250s, SL270s, SL165z 	<ul style="list-style-type: none"> • InfiniBand • Gigabit Ethernet
HP ProLiant Moonshot server	X86_64	<ul style="list-style-type: none"> • RHEL, SLES, Ubuntu 	<ul style="list-style-type: none"> • Moonshot cartridges 	<ul style="list-style-type: none"> • Wolff 1 switch integrated in Moonshot chassis

HP Insight CMU also supports kick-starting compute nodes, as well as deploying Linux in a diskless environment.

The HP Insight CMU GUI displays real-time monitoring of selected metrics from some or all nodes in an intuitive and uniquely scalable design. This capability allows the system manager to view the status of the entire cluster at a glance or to focus on a group of nodes and change the displayed metrics as necessary. A new feature, TimeView, renders the metric display in 3D, providing a time dimension to see cluster behavior across multiple metrics in parallel.

HP Insight CMU provides a rich set of preconfigured metrics and the flexibility to add customized metrics. System administrators can also customize the grouping of nodes to enhance monitoring. HP Insight CMU makes available a script that integrates with various batch schedulers to dynamically create and delete groups of nodes in HP Insight CMU that correspond with the workloads currently active on the cluster. With the HP Insight CMU extended monitoring interface, sensor data can be collected directly from the HP iLO Management Engine—for example, from the HP Agentless Management software.

Cloning feature

HP Insight CMU offers the ability to clone the disk contents of one node to other cluster nodes. This feature reduces the complexity of scale and helps enable a consistent configuration of compute nodes within the cluster.

A cluster is often split into logical groups, which may be owned by different parts of the IT organization. HP Insight CMU can manage these logical groups by associating one or more disk images with each group. Once one node in a group has been installed and configured, HP Insight CMU builds a compressed image (called a “clone image”) of this master disk. The clone image is then ready to be propagated to other members of the group using the HP Insight CMU cloning mechanism.


Cloning chronology

A node that is ready for cloning is first rebooted to free the local disk of any operating system activity. Then the diskless HP Insight CMU imaging environment is loaded on the node, and a backup image of the system disk is captured and stored.

Once a successful backup has occurred, the cloning process can be invoked. Using a tree propagation approach, this process is divided in two steps to protect the network topology.

In the first step, one node from each network entity is given a copy of the backup image.

Once a node has received a completed image, the node tries to download the image to another node within the group. HP Insight CMU manages the group node list waiting for an image. This algorithm parallelizes the propagation process and takes advantage of the available network bandwidth. This high-performance tree propagation algorithm can quickly clone a full operating system and its applications. The propagation mechanism efficiently uses network bandwidth and is specifically designed to capitalize on switched network configurations.

Figure 3. HP Insight CMU TimeView display

Monitoring feature

The monitoring feature was designed to synthesize environmental, performance, and administration information from the cluster for the IT manager. In this way, HP Insight CMU helps decrease total cost of solution ownership. Visibility to the cluster is presented to the administrator in a single system view, as opposed to a compilation of data collected on the compute engines distributed over the cluster. The standard HP Insight CMU interface renders a circular “flower” for each metric, with each flower petal reporting node-specific information. Examples of metrics include CPU and memory utilization, GPU load, and temperature. The new TimeView display renders “tubes” for each user-selected metric. The “tube” is a 3D display of those circular flowers over time, which can also be used to view histories.

Monitoring, similar to all other features of the HP Insight CMU, is designed for scalability. No additional resources are needed to ensure high-performance, highly efficient monitoring—even as a cluster grows to include thousands of nodes.

The primary advantage of HP Insight CMU’s monitoring GUI is to provide a single system view of the cluster. The monitored data of all individual compute nodes displays in a single pie view, providing the administrator with a comprehensive picture of the entire cluster’s behavior.

Comprehensive support and services

Available factory installed, HP Insight CMU software is a featured cluster management option for HP Cluster Platforms. HP Insight CMU software is supported worldwide by the service professionals at HP. HP Insight CMU software can be complemented by a full range of services, from design to factory preinstallation and configuration to acceptance and training. A public forum for [HP Insight CMU](#) users is also hosted under the heading of Servers in the HP Community home website.

Learn more at
hp.com/go/cmu

Sign up for updates
hp.com/go/getupdated



Share with colleagues



Rate this document

© Copyright 2010, 2012–2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

AMD and Opteron are trademarks of Advanced Micro Devices, Inc. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Windows is a U.S. registered trademark of Microsoft Corporation. Java is a registered trademark of Oracle and/or its affiliates.

4AA1-5259ENW, August 2013, Rev. 4

